


Using Sigma-Delta Converters—Part 1

Q: I'd like to use sigma-delta A/Ds but have some questions because they seem markedly different from what I've been using. To start with, what issues do I need to consider when designing my antialiasing filter?

A: A major benefit of oversampling converters is that the filtering required to prevent aliases can be quite simple. To understand why this is the case and what the filter constraints are, let's look at the basic digital signal processing that takes place in such a converter. For the purpose of anti-alias filter design we can think of a sigma-delta converter as a conventional high-resolution converter, sampling at a rate much faster than the Nyquist sampling rate, followed by a digital decimator/filter. That the input into the digital decimator is 1-bit serial with a noise-shaping transfer function doesn't matter.

The input signal is sampled at F_{ms} , the modulator input sampling rate, which is much faster than twice the maximum input signal frequency (the Nyquist rate). The figure shows what the frequency response of a decimation filter may look like; frequency components between f_b and $F_{ms} - f_b$ are greatly attenuated. Thus, the digital filter can be used to filter out all energy from the converter within $[0, F_{ms} - f_b]$ that does not fall within the bandwidth of interest $[0, f_b]$. However, the converter can not distinguish between signals appearing at the input that are in the range $[0, \pm f_b]$ and those in the ranges, $[kF_{ms} \pm f_b]$, where k is an integer. Any signals (or noise) in those ranges get aliased down to the bandwidth of interest $[0, f_b]$ via the sampling process; the decimation filter, which works only on the digitized samples, cannot be of any help attenuating these signals.



Thus it is the input noise energy in these bands $[kF_{ms} \pm f_b]$ that must be removed by the antialiasing filter before the input signal is sampled by the converter.

Q: So if I were to use the AD1877 (available in Spring, 1994), which has a dynamic range of 90 dB, the anti aliasing filter will need attenuation well above 90 dB at $F_{ms} - f_b$ (≈ 3 MHz)?

A: Not quite. You are assuming that the A/D has full-scale input at frequencies close to the modulator sampling rate; this is simply not the case in most systems. The only signal input of concern for aliasing is normally just noise from sensors and circuitry preceding the converter. The noise is usually low enough for a simple RC filter to suffice as an anti-alias filter.

Q: How do I make sure that a one-pole RC filter will suffice for my application—and establish the time constant of the filter?

A: Your application will typically specify a maximum allowable attenuation of an input signal that falls within the bandwidth of interest. This in turn puts a minimum on the -3-dB point of the RC filter. Let's take a look at an example using the AD1877 to illustrate this point further and to show how one might verify that a single-pole filter will provide enough filtering.

Let's assume that we have an application where the bandwidth of interest is 0 to 20 kHz, and signals in this range must not be attenuated more than 0.1 dB, or a ratio of 0.9886 [$\text{dB} = 20 \log_{10}(\text{ratio})$ for voltage and $10 \log_{10}(\text{ratio})$ for power]. From the formula for attenuation of a single-pole filter,

$$\text{ratio} = \frac{1}{\sqrt{1 + (2\pi fRC)^2}} > 0.99 \text{ at } f = 20 \text{ kHz}$$

$$RC \leq \sqrt{\frac{1 - (\text{ratio})^2}{(2\pi f)^2 (\text{ratio})^2}} \approx 1.21 \times 10^{-6} \text{ s}$$

Choosing $RC = 1.0 \mu\text{s}$, to allow for component tolerances, the -3-dB frequency will be 159 kHz. We can now calculate the attenuation the filter will provide in the frequency bands, $kF_{ms} \pm f_b$, that alias down to the baseband. Assuming that the AD1877 has a modulator sampling rate of 3.072 MHz (and output sampling rate of 48 kHz), the first frequency band occurs at 3.052 MHz to 3.092 MHz. The attenuation of the RC filter at these frequencies is approximately 25.7 dB (about 0.052) over the whole band. Over the second band (6.124 MHz to 6.164 MHz), the attenuation is 31.8 dB (0.026). We know

that the noise in these two bands (and all higher bands up the scale) that escapes through the filter to the A/D input will be aliased down to the baseband and get added as root sum-of-the-squares (rss) of their rms values, i.e., $\sqrt{n_1^2 + n_2^2 + \dots + n_i^2}$. For values given in dB, the formulas shown the Appendix can provide results directly in dB, avoiding the intermediate step of computing the ratios.

For white noise, the noise spectral density is constant as a function of frequency, and each frequency range has the same bandwidth, so each band contributes an equal amount of noise to the input of the filter. We can therefore find the effective attenuation of the RC filter by adding the attenuation of the different frequency bands in rss fashion. The noise contribution from the first two bands, for example, is the same as the contribution from a single frequency band with attenuation of $\sqrt{0.052^2 + 0.026^2} = 0.058$, or 24.7 dB, compared with 25.7 dB for the first band. How many bands do we need to consider when calculating the total aliased noise? For this case, the rss sums of the first 3, 4, 5, and 6 bands are, respectively, -24.2, -24.0, -23.9, -23.8 dB. The first band is therefore quite dominant; its attenuation is within 2 dB of the attenuation for all bands. It is usually sufficient to take only the first band into account unless the noise is exceptionally large or has a non-white spectrum; in addition, the A/D itself, though fast, has limited bandwidth; it tends to reject high-order bands.

Now that the attenuation is in hand, we can consider the noise magnitude itself: Let's be conservative (by about 50%) and take the effective filter attenuation to be 20 dB (i.e., 0.1 V/V). To be able to calculate the maximum allowed noise spectral density when using a single pole filter, an estimate should be made of the maximum performance degradation that aliased noise can contribute. From the dynamic specs of the AD1877 we find that the total noise power internal to the converter is 90 dB below (32 ppm of) full-scale input. If the whole system is to be within, say, 0.5 dB of this spec, the total aliased noise power can't exceed the rss difference between -90 dB and -89.5 dB or -99.1 dB (11.1×10^{-6}). Using this information, and the fact that the input scale of the AD1877 is 3 V p-p, we find that aliased noise must not exceed $3/(2\sqrt{2}) \text{ V} \times 11.1 \times 10^{-6} = 11.8 \mu\text{V}$. If all this noise were assumed lumped in a single aliased band, and noting that rms noise = noise spectral density $\times \sqrt{\text{BW}}$,

$$\text{N.S.D.} < \frac{11.8 \mu\text{V}}{\sqrt{3.092 \text{ MHz} \times 3.052 \text{ MHz}}} = 59 \text{ nV}/\sqrt{\text{Hz}}$$

This is the maximum post-filter spectral density allowed. To find the maximum prefilter spectral density (MPSD), with the effective filter attenuation of 20 dB (i.e., $\times = 10$) established previously, M.P.S.D. = $10 \times 59 \text{ nV}/\sqrt{\text{Hz}} = 0.59 \mu\text{V}/\sqrt{\text{Hz}}$.

Clearly your system has to be pretty noisy in the 3-6-9-12-MHz regions in order for a simple RC filter not to suffice; however, as always, one must be careful of ambient rf pickup.



Q: As I understand it, the noise floor of sigma-delta converters may exhibit some irregularities. Any thoughts on that?

A: Most sigma-delta converters exhibit some spikes in the noise floor, called *idle tones*. In general, these spikes have low energy, not enough to substantially affect the S/N of the converter.

Despite that, however, many applications cannot tolerate spikes in the frequency spectrum that extend much beyond the white noise floor. In audio applications, the human ear, for example, does an excellent job of detecting tones in the absence of large input signals even though the tones are well below the integrated (0-20-kHz) noise of the system.

There are two sources of idle tones. Their most common cause is voltage-reference modulation. To understand this mechanism a basic understanding of sigma-delta converters is needed. Here is a one minute crash course on sigma-delta converters (to probe further please consult).[1]

As the block diagram shows, a basic sigma-delta A/D converter consists of an oversampling modulator, followed by a digital filter and a decimator. The modulator output swings between two states (high and low, or 0 and 1, or +1 and -1), and the average output is proportional to the magnitude of the input signal. Since the modulator output always swings full-scale (1 bit), it will have large quantization errors. The modulator, however, is constructed so as to confine most of the quantization noise to the portion of the spectrum beyond f_b , the bandwidth of interest.




As shown, the spectral "sticks" (single frequencies) at f_i and $F_{ms} - f_i$ correspond to an input signal, while the shaded area shows how the quantization noise has been pushed (shaped) beyond the bandwidth of interest, f_b .

The digital filter, which is often an n -tap FIR filter, takes the high-speed low-resolution (1-bit) modulator output and performs a weighted average of n modulator outputs in a manner dictated by the desired filter characteristics. The output of the filter is a high-resolution word, which becomes the A/D output. The digital filter is designed to filter out "everything" between f_b and $F_{ms} - f_b$, where F_{ms} is the sampling rate of the modulator. Cleaning out all the noise in between f_b and $F_{ms} - f_b$ makes it possible to reduce the sampling rate to values between F_{ms} and $2f_b$ without causing any spectra to overlap (i.e., aliasing).


Conceptually, reducing the sample rate, i.e., *decimation*, can be thought of as only sending every d th digital filter output to the A/D output, where d is the decimation factor. This will bring the spectral images close together, as shown in the figure, which

makes the output look like an output from a non oversampled converter. The upper figure shows the output of the modulator after digital filtering but prior to decimation. The lower figure shows the spectral output after decimation—the final A/D output.

In real converters, digital filtering and decimation are intimately combined for economy in design and manufacture. Thus, the terms “digital filter” and “decimator” are used interchangeably to describe the digital circuitry processing the modulator output to produce the output of the converter.



O.K., now back to “idle tones”. Let’s start by looking at the output of the modulator when a dc signal is applied to the input. For an exact mid-scale dc input level, the output of the modulator is equally likely to be high (1) or low (0), in other words, the pulse density is 0.5, very likely to result in bitstream patterns like 010101. These regular patterns mean that the output spectrum will have a spike at $F_{ms}/2$ (upper figure). If the dc input now moves somewhat off midscale, the modulator output bit pattern will change accordingly. The spectrum of the modulator output will now show spikes at $F_{ms}/2 - \partial F$ and $F_{ms}/2 + \partial F$, with ∂F proportional to the dc change from midscale (lower figure).



With effective digital filtering, how can such tones possibly find their way down to baseband? The answer is via the *voltage reference*. The digital output is a measure of the ratio of the analog input to the voltage reference. An $x\%$ change in the magnitude of the voltage reference will result in a $-x\%$ change in the magnitude of the digital output word. Voltage-reference change will, in effect, amplitude modulate the A/D output. Now, we have clocks internal to the converter, and possibly also externally, running at $F_{ms}/2$. If small amounts of these clock pulses get coupled onto the voltage reference line, they will change it slightly and, in effect, modulate the tones at $F_{ms}/2 - \partial F$ and $F_{ms}/2 + \partial F$. One of the difference frequencies created by this modulation is at ∂F , and it is clearly in the bandwidth of interest. Nonlinearities may also create tones at multiples of ∂F .

Q: From your explanation it seems that if I apply an ac signal to the converter I do not have to worry about idle tones?

A: Well, any ac signal generally has a dc component associated with it, which will have to be represented by the modulator output, so the explanation above still applies. But if the total dc input offset (i.e., internal converter offset plus external offsets) in your system is exactly 0, the tones will be at dc (0 Hz).

There is another source of idle tones in lower-order (<3rd-order) modulators. The order of the modulator (number of integrations) is a measure of how much quantization-noise shaping takes place. Second-order modulators can actually exhibit bit patterns that show up directly in the baseband, even if voltage-reference modulation is not occurring. This is one of the reasons why sigma-delta converters from Analog Devices that are designed for ac applications use higher-order (≥ 3) sigma-delta modulators.


Q: So what can I do to minimize the chances of idle tones interfering with my A/D conversion?

A: Follow the layout recommendations and bypassing schemes recommended by the manufacturer of the converter. This applies not only to the voltage reference, but to power supplies and grounding as well. It is the manufacturer’s responsibility to minimize the voltage-reference corruption that takes place inside the converter, but it is up to the system designer to minimize the external coupling. By following those guidelines, the user should be able to reduce the coupling to a negligible level. If, despite the proper design precautions, idle tones are still an issue, there is yet another option that can be pursued. As I explained previously, frequency of the idle tones is a function of the dc input. This opens up the possibility of introducing enough dc offset on the A/D input to move the idle tones out of the bandwidth of interest to where they will be filtered out by the decimation filter. If the user does not want the dc offset to propagate through the system it can be subtracted out by the processor that handles the data from the A/D.

Q: What kind of a load does the input of sigma-delta converters present to my signal conditioning circuitry?

A: It depends on the converter. Some sigma-delta converters have a buffer at the input, in which case the input impedance is very high and loading is negligible. But in many cases the input is

connected directly to the modulator of the converter. A switched-capacitor sigma-delta modulator will have a simplified equivalent circuit like that shown in the figure.



Switches S1 and S2 are controlled by the two phases of a clock to produce alternating closures. While S1 is closed, the input capacitor samples the input voltage. When S1 is opened, S2 is closed and the charge on C is dumped into the integrator, thus discharging the capacitor. The input impedance can be computed by calculating the average charge that gets drawn by C from the external circuitry. It can be shown that if C is allowed to fully charge up to the input voltage before S1 is opened that the average current into the input is the same as if there were a resistor of $1/(F_{sw}C)$ ohms connected between the input and ground, where F_{sw} is the rate at which the input capacitor is sampling the input voltage. F_{sw} is directly proportional to the frequency of the clock applied to the converter. This means that the input impedance is inversely proportional to the converter output sample rate.

Sometimes other factors, such as gain, can influence the input impedance. This is the case for the 16/24-bit AD771x family of signal conditioning A/Ds. The inputs of these converters can be programmed for gains of 1 to 128 V/V. The gain is adjusted using a patented technique that effectively increases F_{sw} (but keeps the converter output sample rate constant) and combines the charges from multiple samples. The input impedance of these converters is, for example, 2.3 M Ω when the device's external clock is 10 MHz and the input gain is 1. With input gain of 8, the input impedance is reduced to 288 k Ω .


These impedances, as noted earlier, represent the average current flow into or out of the converters. However, they are not the impedances to consider when determining the maximum allowable output impedance of the A/D driver circuitry. Instead, one needs to consider the charging time of the capacitor, C, when S1 is closed. For dc applications the driver circuit impedance has only to be low enough so that the capacitor, C, will be charged to a value within the required accuracy before S1 is opened. The impedance will be a function of how long S1 is closed (proportional to the sampling rate), the capacitance, C and C_{EXT} in parallel with the input (unless $C_{EXT} \gg C$). The table shows allowable values of external series resistance with $f_{CLKIN} = 10$ MHz which will avoid gain error of 1 LSB of 20 bits—for various values of gain and external capacitance on the AD7710.

Typical External Series Resistance Which Will Not Introduce 20-Bit Gain Error


Gain	External Capacitance (pF)					
	0	50	100	100	500	5000
1	145 k Ω	34.5 k Ω	20.4 k Ω	5.2 k Ω	2.8 k Ω	700 Ω
2	70.5 k Ω	16.9 k Ω	10 k Ω	2.5 k Ω	1.4 k Ω	350 Ω
4	31.8 k Ω	8.0 k Ω	4.8 k Ω	1.2 k Ω	670 Ω	170 Ω
8–128	13.4 k Ω	3.6 k Ω	2.2 k Ω	550 Ω	300 Ω	80 Ω

For ac applications, such as audio, where the modulator sample rate is around 3 MHz for 64 \times oversampling, the input capacitor voltage may not have enough time to settle within the accuracy indicated by the resolution of the converter before the capacitor is switched to discharging. It actually turns out that as long as the input capacitor charging follows the exponential curve of RC circuits, only the gain accuracy suffers if the input capacitor is switched away too early.

The requirement of exponential charging means that an op amp can not drive the switched capacitor input directly. When a capacitive load is switched onto the output of an op amp, the amplitude will momentarily drop. The op amp will try to correct the situation and in the process hits its slew rate limit (non linear response), which can cause the output to ring excessively. To remedy the situation, an RC filter with a short time constant can be interposed between the amplifier and the A/D input as shown in the figure. The (low) resistance isolates the amplifier from the switched capacitor, and the capacitance between the input and ground supplies or sinks most of the charge needed to charge up the switched capacitor. This ensures that the op amp will never see the transient nature of the load. This additional filter can also provide antialiasing.



For converters that have a differential input, a differential version of this circuit may be used, as shown in the figure below. Since one input is positive with respect to ground while the other is negative, one input (the negative one) needs to be supplied negative charge while the other needs to get rid of negative charge when the input capacitors are switched on line. Connecting a capacitor between the two inputs enables most of the charge that is needed by one input to be effectively supplied by the other input. This minimizes undesirable charge transfers to and from the analog ground.



APPENDIX

RSS addition of logarithmic quantities: The root-square sum of two rms signals, S_1 and S_2 , has an rms value of $\sqrt{S_1^2 + S_2^2}$. One often needs to calculate the rss sum of two numbers that are expressed in dB relative a given reference. To do this one has to take the antilogs, perform the rss addition, then convert the result back to dB. These three operations can be combined into one convenient formula: If D_1 and D_2 are ratios expressed in dB [negative or positive] their sum, expressed in dB, is

$$10 \log_{10} (10^{D_1/10} + 10^{D_2/10})$$

Similarly, to find the difference between two rms quantities,

$$x = \sqrt{S_2^2 - S_1^2}$$

the result, x , expressed in dB, is

$$10 \log_{10} (10^{D_2/10} - 10^{D_1/10})$$

References (not available from Analog Devices):

¹*Oversampling Delta-Sigma Data Converters—Theory, Design, and Simulation*, edited by J.C. Candy and G.C. Temes, IEEE Press, Piscataway, NJ, 1991.

²J. Vanderkooy and S.P. Lipshitz, "Resolution Below the Least Significant Bit in Digital Systems with Dither," *J. Audio Eng. Soc.*, vol. 32, pp. 106-113 (1984 Mar.); correction *ibid.*, p.889 (1984 Nov.).

³A.H. Bowker and G.J. Lieberman, *Engineering Statistics*, Prentice Hall, Englewood Cliffs, NJ, 1972.